

The impact and interplay of long and short branches on phylogenetic information content

Iain Martyn and Mike Steel*

*Biomathematics Research Centre,
University of Canterbury, Christchurch, New Zealand*

Abstract

In molecular systematics, evolutionary trees are reconstructed from sequences at the tips under simple models of site substitution. A central question is how much sequence data is required to reconstruct a tree accurately? The answer depends on the lengths of the branches (edges) of the tree, with very short and very long edges requiring long sequences for accurate tree inference, particularly when these branch lengths are arranged in certain ways. For four-taxon trees, the sequence length question has been investigated for the case of a rapid speciation event in the distant past. Here, we generalize results from this earlier study, and show that the same sequence length requirement holds even when the speciation event is recent, provided that at least one of the four taxa is distantly related to the others. However, this equivalence disappears if a molecular clock applies, since the length of the long outgroup edge becomes largely irrelevant in the estimation of the tree topology for a recent divergence. We also discuss briefly some extensions of these results to models in which substitution rates vary across sites, and to settings where more than four taxa are involved.

Keywords: phylogenetic tree, sequences, Markov model, information content, site saturation

* Corresponding author

Email address: iain.martyn@mail.mcgill.ca, mike.steel@canterbury.ac.nz (Iain Martyn and Mike Steel*)

Preprint submitted to Elsevier

March 8, 2013

1. Background

Phylogenetic methods are founded on the notion that evolutionary relationships can be inferred from sequences that have evolved along with the taxa. It is usually supposed that such sequences evolve according to some continuous-time reversible Markov process, or a mixture of such processes (for further background on phylogenetic inference, the reader is referred to Felsenstein (2004)). Here, we are interested in the question of the sequence length required to accurately estimate a discrete and fundamental parameter of evolutionary history, namely the topology of the underlying evolutionary tree. This question has long been of interest in molecular systematics (see, for example, Saitou and Nei (1986), Churchill et al. (1992) and Lecointre et al. (1994)) and a variety of mathematical approaches have been explored in order to quantify how much ‘phylogenetic information’ sequence data contains (Mossel and Steel (2005), Townsend (2007), Townsend and Leuenberger (2011) and Townsend et al. (2012)). Although the underlying tree topology is rooted, phylogenetic models are generally time-reversible, and so methods based on these models produce trees that are unrooted; accordingly, we will say that a method correctly reconstructs the tree topology if it does so up to the placement of the root.

Amongst unrooted trees, the simplest phylogenetic problem involves a set of four taxa, for which there are just three resolved binary tree topologies and one ‘star tree’. Fischer and Steel (2009) investigated the sequence length required to accurately reconstruct a binary four-taxon phylogenetic tree with four long pendant branches, and a short interior edge (Fig. 1(a), 1(a')). This special case is motivated by the scenario in evolutionary biology in which a rapid speciation event in the distant past results in all taxa sitting on ‘long branches’ around a short interior edge of length l_0 . The authors found that the length of sequence needed to reconstruct the correct four-taxon tree with probability $1 - \epsilon$ grows at the rate Ce^{bL}/l_0^2 , where C and b are positive constants and L is the length of the long pendant edges. Notice the impact on the required sequence length of a long branch length (i.e. $e^{bL} \rightarrow \infty$ as $L \rightarrow \infty$) and of a short interior branch (i.e. $\frac{1}{l_0^2} \rightarrow \infty$ as $l_0 \rightarrow 0$), and that these combine multiplicatively in this lower bound (thus, the cumulative effect of

a short branch beside a long one becomes compounded much more than if the interaction was, say, additive). This formally justifies the informal notion that a very short interior edge surrounded by long branches is a particularly challenging phylogenetic problem.

In this paper, we wish to compare this scenario with another that is at least as common in evolutionary biology, namely the setting in which only one of the taxa is distantly related to the others, being a distant ‘outgroup’ taxon (see, for example, Fig. 1(b)). In particular, we ask whether the sequence length requirements are less severe if just one pendant edge is long, rather than all four. We show analytically that essentially the same bound (of the form Ce^{bL}/l_0^2) applies in general.

However, a curious situation develops if one imposes a molecular clock. Doing so does not affect the exponential depending on L of the sequence length requirements for reconstructing the tree in Fig. 1(a’). However, in the case of just one distant outgroup taxon (Fig. 1(b’)), the exponential dependence on L (the long branch) disappears entirely.

Finally, we extend our main result to settings where sites evolve at varying rates and we also indicate how the results apply when the four lineages are replaced by four monophyletic groups of taxa.

2. Preliminaries

We first recall some terminology from phylogenetics. Let $X = \{1, 2, 3, 4\}$ be a set of four taxa, and let T_1, T_2 and T_3 be the three possible unrooted binary trees that have X as their leaf set.

Suppose we have a continuous, stationary, and time-reversible Markov process on a state space G that acts at various intensities on the edge of one of these trees. The *length* of an edge will refer to the expected number of substitutions on that edge. This is the substitution rate on that edge, multiplied by the temporal duration of that edge. In the case where the substitution rate is constant across the tree, we will say that a *molecular clock* applies, but we do not assume this unless otherwise stated. Throughout this paper we will let l_0 be the branch length of the interior edge of any four-taxon tree.




Figure 1: In tree (a), a short interior edge is incident with four long pendant edges, representing a rapid radiation event deep in the past; tree (a') shows the associated unrooted tree. Tree (b) shows a more recent rapid radiation event in which only one of the four incident edges is long, as it joins a distant outgroup. Tree (b') shows the associated unrooted tree.

Let $S = G^X$ be the set of possible assignments of elements of the state space G to the leaf set X ; we will refer to an element of S as a *site pattern*. Now, suppose we generate k site patterns independently according to the same Markov process to form sequences of length k (one sequence for each taxon). It is well known that for any set of (positive) branch lengths on T_i , one can correctly recover the topology T_i from these sequences with a probability of at least $1 - \epsilon$ for sufficiently large values of k (for a discussion of statistically consistent tree reconstruction, see Felsenstein (2004)). Here ‘sufficiently large’ depends not just on ϵ but also on the tree and its associated branch lengths.

As in Fischer and Steel (2009), our arguments rely on the properties of the *Hellinger*

distance (d_H), which is defined as follows: Given a finite set U , the Hellinger distance $d_H(p, q)$ between two probability distributions p and q on U is defined by the equation:

$$(1) \quad d_H^2(p, q) = \sum_{u \in U} (\sqrt{p_u} - \sqrt{q_u})^2 = 2(1 - \sum_{u \in U} \sqrt{p_u q_u}).$$

Hellinger distances are useful to quantify the amount of data required to accurately identify a discrete parameter in a stochastic model. We will describe this below (Lemma 2.1) in a general setting, not specific to phylogenetics.

2.1. Hellinger distance bounds on required sequence length

Let A and U be finite sets, and suppose that each element $a \in A$ defines a probability distribution p_a on U . We will denote the Hellinger distance between p_a and p_b by $d_H(a, b)$, and, by slight abuse of terminology, refer to it as the ‘Hellinger distance between a and b ’.

Suppose an element ξ of A is selected according to some discrete non-zero probability distribution on A . Conditional on $\xi = a$, consider a sequence of k samples of U generated independently in U according to the probability distribution p_a . Let $M : U^k \rightarrow A$ be some method for estimating the element $a \in A$ from a sequence $(u_1, u_2, \dots, u_k) \in U^k$ (here M may be a deterministic function from U^k to A or a process that selects an element of A from each element of U^k according to some probability distribution – the latter case allows ties to be broken randomly).

Let $r_a^{(M, k)}$ denote the probability that the method M correctly identifies the element a that generates the sequence $(u_1, u_2, \dots, u_k) \in U^k$ under the probability distribution p_a . In other words:

$$r_a^{(M, k)} = \mathbb{P}(M(u_1, u_2, \dots, u_k) = a | \xi = a).$$

The following lemma is from Steel and Székely (2002) (Theorem 3.1 and (2.7)).

Lemma 2.1 *Given finite sets U and A , suppose that elements of U are generated i.i.d. by some unknown element $\xi \in A$. Then for any estimation method M that satisfies*

$r_a^{(M,k)} \geq 1 - \epsilon$, for all $a \in A$, we must have

$$k \geq \frac{C_\epsilon}{d^2},$$

where $C_\epsilon = \frac{1}{4}(1 - \frac{|A|}{|A|-1}\epsilon)^2$, and $d = \min\{d_H(a, a') : a, a' \in A; a \neq a'\}$.

In our setting, A will consist of a set of phylogenetic trees on the leaf set $X = \{1, 2, 3, 4\}$ and U will be the set S of assignment of states of the elements of X . We will use the lemma to prove a lower bound on k in the following Section.

3. A general lower bound on the required sequence length

We now present a lower bound for the necessary sequence length k required to reconstruct a tree of the type shown in Fig. 1(b). This lower bound is essentially of the same form as that which applies when all four pendant branches are long – namely, it grows exponentially with the length L of the long branch and in inverse proportion to the square of the short interior branch, and these factors combine multiplicatively.

Theorem 3.1 *Consider the three-leaf star tree on the taxon set $\{1, 2, 3\}$ with corresponding branch lengths $l_1, l_2, l_3 \geq \delta > 0$. Suppose that a fourth taxon is attached by a branch of length $L > 0$ to one of the three branches at a distance $l_0 \in (0, \delta)$ from the interior node. Generate k i.i.d. site patterns at the tips of the resulting four-taxon tree under a finite-state, stationary and irreducible Markov process. Then, any method that is able to correctly identify with probability at least $1 - \epsilon$ which branch the fourth branch is grafted onto requires:*

$$(2) \quad k \geq Ce^{bL}/l_0^2$$

where C is a constant (independent of l_0 and L) that depends on ϵ, δ and the rate parameters of the Markov process.

Moreover, some methods achieve this accuracy using sequences with a length that is no more than a constant times e^{bL}/l_0^2 .




Figure 2: The three-taxon and four-taxon trees described in the statement and proof of Theorem 3.1

Proof: To establish Inequality (2) as a lower bound, we first derive an upper bound for the Hellinger distance between T_1 and the tree T^* , formed by grafting the fourth branch directly onto vertex u of the three-taxon star tree, as shown in Fig. 2. By the triangle inequality, we have:

$$(3) \quad d_H(T_1, T_2) \leq d_H(T_1, T^*) + d_H(T^*, T_2).$$

Most of the proof is devoted to establishing the following inequality (for a constant $B = B(\delta)$):

$$(4) \quad d_H^2(T_1, T^*) \leq B l_0^2 e^{-bL}.$$

To establish Inequality (4), let p_s and p_s^* denote the probability of generating the site pattern s on T_1 and T^* , respectively. Let N denote the number of substitutions occurring along the edge between u and v (the interior edge of T_1), and let $\tau = \mathbb{P}(N > 0)$. Further, let Q_s (respectively Q_s^*) denote the conditional probability of generating pattern s on T (respectively, on T^*) given that $N = 0$. Similarly, let P_s (respectively P_s^*) denote the conditional probability of generating pattern s on T (respectively, on T^*) given $N > 0$. Then, by the law of total probability, we can write p_s and p_s^* as:

$$(5) \quad \begin{aligned} p_s &= (1 - \tau) \cdot Q_s + \tau \cdot P_s, \\ p_s^* &= (1 - \tau) \cdot Q_s^* + \tau \cdot P_s^*. \end{aligned}$$

Let $D_s = P_s - P_s^*$. Then, since $Q_s = Q_s^*$, we have:

$$(6) \quad p_s - p_s^* = \tau(P_s - P_s^*) = \tau D_s.$$

Now, the Hellinger distance between T_1 and T^* (on site patterns) is:

$$(7) \quad d_H^2(T_1, T^*) = 2\left(1 - \sum_{s \in S} \sqrt{p_s p_s^*}\right).$$

Following the approach of Fischer and Steel (2009) (Lemma 5.1), substituting $p_s^* = p_s - \tau D_s$ (from Eqn. (6)) into (7) gives:

$$d_H^2(T_1, T^*) = 2 \left(1 - \sum_{s \in S} p_s \sqrt{1 - \frac{\tau D_s}{p_s}} \right)$$

and then the application of the inequality $\sqrt{1+y} \geq 1 + y/2 - y^2/2$ for $y \geq -1$ leads to the following inequality:

$$(8) \quad d_H^2(T_1, T^*) \leq \tau^2 \cdot \sum_{s \in S} \frac{D_s^2}{p_s}.$$

Now, $\tau = \mathbb{P}(N > 0) \leq \mathbb{E}(N)$, so we have $\tau \leq l_0$, and thus we can replace τ in (8) by l_0 to obtain:

$$(9) \quad d_H^2(T_1, T^*) \leq l_0^2 \cdot \sum_{s \in S} \frac{D_s^2}{p_s}.$$

Referring again to Fig. 2, let χ_u be the character state on vertex u , let χ_v be the character state on vertex v , and let χ_1, \dots, χ_4 be the respective character states on the leaf set $\{1, 2, 3, 4\}$, and let $p(s, s')$ denote the conditional probability $\mathbb{P}(\chi_u = s', \chi_v = s'' | N > 0)$.

Then we have:

$$P_s = \sum_{s', s'' \in S} p(s, s') \mathbb{P}(\chi_1 = s^1 | s') \mathbb{P}(\chi_2 = s^2 | s') \mathbb{P}(\chi_3 = s^3 | s'') \mathbb{P}(\chi_4 = s^4 | s''),$$

and

$$P_s^* = \sum_{s', s'' \in S} p(s, s') \mathbb{P}(\chi_1 = s^1 | s') \mathbb{P}(\chi_2 = s^2 | s') \mathbb{P}(\chi_3 = s^3 | s'') \mathbb{P}(\chi_4 = s^4 | s'),$$

where $\mathbb{P}(\chi_i = s^i | x)$ (for $x = s'$ or s'') is probability of generating leaf state s_i at leaf i conditional on state x at the vertex of the tree adjacent to leaf i .

Thus, $|D_s| = |P_s - P_s^*|$ is bounded above as follows:

$$(10) \quad |D_s| \leq \sum_{s', s'' \in S} p(s, s') \mathbb{P}(\chi_1 = s^1 | s') \mathbb{P}(\chi_2 = s^2 | s') \mathbb{P}(\chi_3 = s^3 | s'') \cdot \left| \left(\mathbb{P}(\chi_4 = s^4 | s'') - \mathbb{P}(\chi_4 = s^4 | s') \right) \right|$$

We now invoke the property that any irreducible Markov process converges to its stationary distribution at an exponential rate regardless of its starting state (*cf.* Theorem

8.3 of Rozanov (1969)). Specifically, if Y_t is the state of such a process when it is run for duration t then, for any state s with the equilibrium frequency $\pi(s)$, and any second state α , we have:

$$(11) \quad |\mathbb{P}(Y_t = s | Y_0 = \alpha) - \pi(s)| \leq Ae^{-at},$$

where A and a depend only on the rate parameters of the Markov process. Using the triangle inequality, Inequality (11) gives:

$$(12) \quad |\mathbb{P}(\chi_4 = s^4 | s') - \mathbb{P}(\chi_4 = s^4 | s'')| \leq 2Ae^{-aL}.$$

Substituting Eqn. (12) into Eqn. (10), we obtain:

$$(13) \quad \frac{D_s^2}{p_s} \leq \frac{4A^2e^{-2aL}}{p_s} \left(\sum_{s', s'' \in S} p(s, s') \mathbb{P}(\chi_1 = s^1 | s') \mathbb{P}(\chi_2 = s^2 | s') \mathbb{P}(\chi_3 = s^3 | s'') \right)^2$$

and since the term in brackets is bounded above by $1 (= \sum_{s', s'' \in S} p(s, s'))$, we obtain:

$$(14) \quad \frac{D_s^2}{p_s} \leq \frac{4A^2e^{-2aL}}{p_s}.$$

Note also that, since $l_1, l_2, l_3 \geq \delta > 0$ and $L > 0$, and the Markov process is irreducible, there is some positive $\rho = \rho(\delta)$ such that $p_s \geq \rho$. We can thus further reduce Inequality (13) to:

$$(15) \quad \frac{D_s^2}{p_s} \leq \frac{4A^2e^{-2aL}}{\rho}.$$

Substituting Eqn. (15) into Eqn. (9) now furnishes the promised justification of Eqn. (4), upon taking $b = 2a$ and $B = B(\delta) = 4A^2/\rho(\delta)$. By symmetry, Eqn. (4) gives us the same upper bound on $d_H^2(T_2, T^*)$ as for $d_H^2(T_1, T^*)$. We then have, by the triangle

inequality:

$$(16) \quad d_H^2(T_1, T_2) \leq (d_H(T_1, T^*) + d_H(T^*, T_2))^2 < 4Bl_0^2 e^{-bL}.$$

The first part of Theorem 3.1 follows from Lemma 2.1 by taking $A = \{T_1, T_2\}$ (so that $C_\epsilon = \frac{1}{4}(1 - 2\epsilon)^2$) and then setting $C = \frac{C_\epsilon}{4B(\delta)}$.

Finally, the last claim in Theorem 3.1 (that e^{bL}/l_0^2 is an upper bound on the required sequence length, up to a constant multiplicative factor) is provided by Theorem 14 of Erdős et al. (1999).

□

3.1. Imposing a (relaxed) molecular clock

When a molecular clock is imposed, there is an interesting shift in the sequence length requirements for accurate tree reconstruction. Although we have seen that the two scenarios in Fig. 1 lead to the same type of lower-bound dependence of sequence length on l_0 and L , namely $\exp(cL)/l_0^2$, if we impose a molecular clock, then this equivalence disappears. More precisely, it is clear (from Fischer and Steel (2009)) that the term $\exp(cL)/l_0^2$ remains for the deep divergence set-up of our Fig.1(a), but for the recent divergence event shown in Fig. 1(b) we will show that the length of the long edge L is largely irrelevant.

We need to stress here how this result should be interpreted. We are not claiming that if a clock applies in the tree that generates the data, then every consistent model-based method, such as maximum likelihood, will be immune to the effect of a long branch to an outgroup. It will not be so immune if, in the model assumed in the maximum likelihood analysis, a molecular clock is not imposed. We are merely claiming that certain methods (such as agglomerative clustering, or MLE with a clock) can be immune to a long branch if a clock assumption applies.

We formalize this by a result, in which the full strength of the molecular clock condition can be relaxed slightly. Note that for the tree in Figure 1b, then under a strict

molecular clock the branch lengths (as indicated in Fig. 2) must satisfy:

$$l_1 = l_2; \ l_3 = l_1 + l_0 \leq L.$$

We relax this slightly by requiring only that:

$$(17) \quad \min\{l_3, L\} \geq \max\{l_1, l_2\} + l_0.$$

Theorem 3.2 *Consider the tree in Fig. 1(b) and suppose that the branch lengths (as indicated in Fig. 2) satisfy the relaxed clock condition described in (17). Let k sites evolve i.i.d along this tree under a finite-state, stationary and reversible Markov process. Then the placement of the branch leading to taxon 4 can be determined correctly with probability at least $1 - \epsilon$ provided that:*

$$(18) \quad k \geq B/(1 - e^{-\lambda l_0})^2,$$

where B depends just on l_1 , the model and ϵ , and where λ is a constant determined by the model. In particular, this bound is independent of the length L of the long branch to the outgroup taxon 4.

Proof: Consider the following simple reconstruction method (this is, essentially, unrooted UPGMA on four taxa). Let $s(x, y)$ denote the proportion of sequence sites for which taxa x and y have the same state. Select the two taxa that maximize $s(x, y)$ and return the (unrooted) quartet tree in which x and y form a cherry. Let $e(x, y)$ be the expected value of $s(x, y)$. If Y_t ($t \geq 0$), denotes the Markov process described in the statement of Theorem 3.2 then we have:

$$e(x, y) = \mathbb{E}[s(x, y)] = \sum_i \pi_i \mathbb{P}(Y_t = i | Y_0 = i),$$

where in this equation that value $t = t_{xy}$ refers to the branch length distance between

taxon x and y . By the spectral representation of reversible continuous-time Markov processes (see e.g. Chapter 3, Eqn (40) of Aldous and Fill (2010)) we have, for any state i :

$$\mathbb{P}(Y_t = i | Y_0 = i) = \pi_i + \sum_{m \geq 2} u_{im}^2 e^{-\lambda_m t},$$

where π_i is the equilibrium frequency of state i , $\lambda_m \geq 0$ are the eigenvalues of the rate matrix multiplied by -1, and the u_{jm} values are real coefficients related to the eigenvalues of the rate matrix. The λ_m values can be ordered $0 = \lambda_1 < \lambda_2 \leq \lambda_3 \leq \dots$. Consequently, $e(x, y) = \sum_i \pi_i^2 + \sum_{m \geq 2} c_m e^{-\lambda_m t}$, where $c_m = \sum_i \pi_i u_{im}^2 > 0$, and so:

$$e(x, y) - e(x', y') = \sum_{m \geq 2} c_m (e^{-\lambda_m t_{xy}} - e^{-\lambda_m t_{x'y'}}).$$

Since the coefficient c_2 is strictly positive, if $t_{x'y'} - t_{xy} \geq l_0$ we can write:

$$(19) \quad e(x, y) - e(x', y') \geq c_2 e^{-\lambda t_{xy}} (1 - e^{-\lambda l_0}) > 0,$$

where, for convenience, we let λ denote λ_2 . Notice that $t_{12} = l_1 + l_2$, and $t_{13} = l_1 + l_3$ and $t_{23} = l_2 + l_3$ and so, by the relaxed clock condition (17) we have $t_{23} - t_{12} \geq l_0$ and $t_{13} - t_{12} \geq l_0$. Thus (19) holds for $(x, y) = (1, 2)$ and $(x', y') = (1, 3), (2, 3)$.

Next, if we $X_{12;3} = s(1, 2) - s(1, 3)$, then observe that:

$$(20) \quad \mathbb{P}(s(1, 2) < s(1, 3)) = \mathbb{P}(X_{12;3} < 0) = \mathbb{P}(X_{12;3} - \mathbb{E}[X_{12;3}] < -\mathbb{E}[X_{12;3}]).$$

In order to exhibit an upper bound this probability, we will apply McDiarmid's inequality (McDiarmid, 1989). First, observe that we can express $s(1, 2) - s(1, 3)$ as a sum of k independent random variables (one for each site), each taking a value of $+1, 0$ or -1 , and this sum has the property that changing any one of these variables (while keeping the others fixed) alters $s(1, 2) - s(1, 3)$ by an additive factor whose absolute value is at most $2/k$. Applying the McDiarmid inequality, noting that: $e(1, 2) - e(1, 3) \geq c_2 e^{-\lambda t_{12}} (1 - e^{-\lambda l_0})$

from Inequality (19), we obtain, from Eqn. (20) that:

$$\mathbb{P}(s(1, 2) < s(1, 3)) \leq \exp(-kc_2^2 e^{-2\lambda t_{12}}(1 - e^{-\lambda t_0})^2/2),$$

and this can be made less or equal to $\epsilon/5$ whenever Inequality (18) is satisfied for $B = \frac{2 \ln(5/\epsilon)}{c_2^2 e^{-2\lambda t_{12}}}$. By symmetry, $\mathbb{P}(s(1, 2) < s(2, 3))$ is also less or equal to $\epsilon/5$ for this value of k . Moreover, by the relaxed clock condition, we also have:

$$\mathbb{P}(s(1, 2) < s(x, 4)) \leq \epsilon/5 \text{ for } x = 1, 2, 3.$$

Thus, with probability at least $1 - \epsilon$, the pair $\{1, 2\}$ will have the strictly largest s -value; consequently, the correct tree topology will be recovered by the method described with probability at least $1 - \epsilon$.

□

4. Further extensions and concluding comments

4.1. Rates across sites

When sites evolve i.i.d. the sequence length required to reconstruct the tree in Fig. 1(a) accurately grows exponentially with the length of L of the long exterior branches; the same holds also for the tree in Fig. 1(b) in the absence of any molecular-clock assumption (Theorem 3.1). We point out that these conclusions need not hold when the sites evolve independently but not identically under a model that allows substitution rates to vary across sites, provided this rate distribution allows arbitrarily small rates, and with appropriate density. Suppose, for example, that site i has rate $r_i = \frac{1}{i}$ for $i = 1, 2, \dots$. Let T' be either an alternative binary tree to T_1 or the unresolved tree (i.e. $T' = T_2$ or T^*), and let $D_H^2(T_1, T')$ be the Hellinger distance between *sequences* of length k generated by T_1 and T' in which the rates at site i of the Markov process is r_i . We claim that for a sequence length that grows at the (polynomial) rate L^5 , the value $D_H^2(T_1, T')$ converges to 2 as L tends to infinity. We first establish this claim and then explain why it implies that

one can reconstruct the generating tree (T_1) from sequences of a length that is polynomial in L .

By a standard equality relating Hellinger distance of sequences of independent samples to the Hellinger distances at each sequence site (easily derived from Eqn. (1)) we have:

$$(21) \quad D_H^2(T_1, T') = 2 \left(1 - \prod_{i=1}^k \left(1 - \frac{1}{2} d_i^2 \right) \right),$$

where d_i is the Hellinger distance between the probability distributions on patterns at site i (and rate r_i) generated by tree T_1 and generated by T' . This applies in either the setting of Figure 1a (four long pendant edges) or Figure 1b (one long pendant edge). Moreover, by definition,

$$(22) \quad d_i^2 \geq (\sqrt{p_i} - \sqrt{p'_i})^2$$

where p_i here refers to the probability of generating a site pattern with leaves 1, 2 in one state (say A) and leaves 3, 4 in a different state (say B) on T_1 at substitution rate λ_i , while p'_i is the corresponding probability for this same site pattern when T_1 is replaced by T' . Notice that this site pattern can be generated by a state change on just one edge of T (the central edge), while on T' at least two pendant edges require state changes. Thus, for suitable constants c, c' , for the tree in Fig. 1(a) we have: $p_i \geq \frac{c}{i}$ and $p'_i \leq c'(\frac{L}{i})^2$; while for the tree in Fig. 1(b) we have $p_i \geq \frac{c}{i}$ and $p'_i \leq c'(\frac{L}{i})(\frac{1}{i})$. Thus, in either case, provided i is sufficiently large, Inequality (22) gives:

$$(23) \quad d_i^2 \geq \left(\frac{\sqrt{c}}{\sqrt{i}} - \frac{\sqrt{c'}L}{i} \right)^2 \geq \frac{d}{i}(1 - o(1)),$$

for a positive constant d , and where $o(1)$ denotes a term that converges to 0 with increasing

i. Then, by combining Eqn. (21) and Eqn. (23), we have:

$$(24) \quad 2 \geq D_H^2(T_1, T') \geq 2 \left(1 - \prod_{i=L^4}^{L^5} \left(1 - \frac{1}{2} d_i^2 \right) \right) \geq 2 \left(1 - \prod_{i=L^4}^{L^5} \left(1 - \frac{d(1 - o(1))}{2i} \right) \right),$$

and straightforward asymptotic analysis of the last term reveals that $D_H^2(T_1, T') \rightarrow 2$ as $L \rightarrow \infty$.

Finally, we invoke an inequality (Theorem 3.2) from Steel and Székely (2002). If $M = \text{MLE}$ (maximum likelihood estimation) then for $A = \{T_1, T'\}$, the probability that MLE correctly reconstructs the generating tree from A is at least $\frac{1}{2} D_H(T_1, T')$ and this converges to 1 as L grows (with k growing at the rate L^5).

4.2. Breaking up long edges by adding more taxa

Based on simulation studies and qualitative understanding, it is received wisdom that long branches are untrustworthy due to the long branch being able to ‘go anywhere’. Hence biologists seek to break up this branch either with more characters or more taxa (see, for example, Felsenstein (2004) or Graybeal (1998)). One could then reconstruct a phylogenetic tree for this ‘extended’ set of taxa and then ignore all but the few taxa one is interested in.

However, as we add more taxa, the number of possible phylogenetic trees grows exponentially, and more data are required to reconstruct a larger tree correctly (this can easily be seen by a purely counting argument). Thus it is not immediately clear whether this strategy has any formal basis for improving accuracy. Here, we show that the sequence length requirements for resolving a four-taxon tree that has one long branch can be exponentially (or even double-exponentially) greater than those of the large tree in certain ideal situations.

To see this, suppose we have one of the types of trees shown in Fig. 1, with one or more long pendant edges of length L . Suppose one can find a set S of N additional taxa so that each edge in the resulting tree has a branch length that lies between fixed values, say l and l' (with $l \leq l_0 \leq l'$). Reconstructing this larger tree accurately requires just

some constant times $\log(N)/l^2$ sites under a two-state symmetric model (see Daskalakis et al. (2011)) provided l' lies below a critical transition value, while reconstructing the four-taxon tree involves a term (e^{cL}) that grows exponentially with the length L of any long edge (by Theorem 3.1).

The significance of this result hinges on the following question: how does $\log(N)$ compare with e^{cL} ? If very short branches are attached at equally spaced intervals along the long pendant branch (or branches), then N grows in a linear relationship with L . In this case, the sequence length required to reconstruct the four-taxon tree is *doubly exponential* in the sequence length required to reconstruct the much larger tree, as L grows (moreover, this does not require the strong technical result from Daskalakis et al. (2011) but a weaker though more generally applicable result from Erdős et al. (1999)).

However, it would be more realistic to constrain the branch lengths in the tree to be approximately clocklike. In that case, N need only be of order 2^{dL} for some constant d ; $\log(N)$ would then be proportional to L and the sequence length required to reconstruct the four-taxon tree would be exponential in the sequence length required to reconstruct the much larger tree as L grows.

In this analysis we are, of course, assuming the most ideal situation, where the taxa are distributed as favorably as possible to allow the large tree to be reconstructed; still, it is interesting to note that this route – constructing a large tree accurately, then ignoring the majority of taxa to consider just the induced phylogenetic relationship between four taxa – can require much shorter sequences lengths to achieve the same accuracy (and this holds for statistically consistent tree reconstruction methods, not just for inconsistent methods that can be ‘misled’ by long branches).

4.3. Extension of Theorem 3.1 to trees with more taxa

Finally, we discuss what happens to our main results concerning four-taxon trees if we replace one or more of the four leaves of the tree by subtrees. Firstly, the lower bound on k given by Theorem 3.1 still applies if the l_i values refer to the lengths of the central three edges. This is because the sequences at the root of the four subtrees screens off the

states of the leaves from the random variable that is the topology T of the central part of the tree (by the Markov property). More formally, consider the following two data sets:

- the sequences Z at the leaves of the tree;
- the sequences Y at the roots of the four subtrees;

Since $T \rightarrow Y \rightarrow Z$ is a Markov chain, the ‘data processing inequality’ (Cover and Thomas (1991)) ensures that $I(T; Z) \leq I(T; Y)$, where I refers to mutual information. In other words, the information that the leaves of the tree tell us about T (the topology of the central part of tree) cannot exceed the information that the ancestral sequences at the roots of those subtrees provide about T (were these known; recall that we only observe sequences at the leaves of the tree). Thus we obtain a conservative lower bound on the required sequence length with these considerations.

However, a tighter bound would presumably take into account how much uncertainty there is in the state at the root of one of the four subtrees, given the states we observe at the leaves of that tree.

To simplify the discussion here, consider just the symmetric two-state model of site substitution. In this case, let p_i denote the probability of accurately inferring the root state of a subtree that stands in place of taxon i from the states at the leaves under maximum likelihood (we assume that the topology and branch lengths of the subtree are known). If l_i is the length of the central branch of T that is incident with the root of this subtree, then the probability of a substitution across the endpoints of this edge is $p'_i = \frac{1}{2}(1 - e^{-2l_i})$.

This suggests the possibility of approximating the sequence length required to resolve a polytomy in a large tree by replacing each of the four incident subtrees by a single taxon, with a net probability of substitution across branch i being set to $p_i \cdot p'_i$. Thus we have replaced a phylogenetic tree with four subtrees by a four-taxon tree, in which the central edge is of the same length but the pendant edges have been ‘lengthened’ to allow for the loss of information that the leaves provide concerning the root state of each

subtree. A natural candidate for this ‘effective branch length’ of branch i would be a value of l for which $p_i \cdot p'_i = \frac{1}{2}(1 - e^{-2l})$; this has the solution: $l = l_i + \frac{1}{2} \log(1/(1 - 2p_i))$. It may be interesting to explore this approach further since the computation and behavior of the expected root-state reconstruction probability (p_i) have been analyzed already by a number of authors (e.g. Evans et al. (2000), and Ma and Zhang (2011)).

5. Acknowledgments

We thank the *Allan Wilson Centre for Molecular Ecology and Evolution* for supporting this work.

References

- Aldous, D., Fill, J., 2010. Reversible Markov chains and random walks on graphs.
- Churchill, G., von Haeseler, A., Navidi, W., 1992. Sample size for a phylogenetic inference. *Mol. Biol. Evol.* 9, 753–769.
- Cover, T., Thomas, J., 1991. *Elements of Information Theory*. Wiley, New York.
- Daskalakis, C., Mossel, E., Roch, S., 2011. Evolutionary trees and the Ising model on the Bethe lattice: a proof of Steel’s conjecture. *Probab. Theor. Relat. Field* 149, 149–189.
- Erdős, P.L., Steel, M.A., Székely, L., Warnow, T., 1999. A few logs suffice to build (almost) all trees (part 2). *Theor. Comput. Sci.* 221, 77–118.
- Evans, W., Kenyon, C., Peres, Y., Schulman, L., 2000. Broadcasting on trees and the Ising model. *Ann. Appl. Probab.* 10, 410–433.
- Felsenstein, J., 2004. *Inferring Phylogenies*. Sinauer Associates, Sunderland, MA.
- Fischer, M., Steel, M., 2009. Sequence length bounds for resolving a deep phylogenetic divergence. *J. Theor. Biol.* 256, 247–252.

- Graybeal, A., 1998. Is it better to add taxa or characters to a difficult phylogenetic problem? *Syst. Biol.* 47, 9–17.
- Lecointre, G., Philippe, H., Van Le, H., Le Guyader, H., 1994. How many nucleotides are required to resolve a phylogenetic problem? The use of a new statistical method applicable to available sequences. *Mol. Phyl. Evol.* 3, 292–309.
- Ma, B., Zhang, L., 2011. Efficient estimation of the accuracy of the maximum likelihood method for ancestral state reconstruction. *J. Combin. Optimization* 21, 409–422.
- McDiarmid, C., 1989. On the method of bounded difference, in: *Surveys in Combinatorics*. Cambridge University Press, Cambridge, pp. 148–188.
- Mossel, E., Steel, M., 2005. How much can evolved characters tell us about the tree that generated them?, in: Gascuel, O. (Ed.), *Mathematics of Evolution and Phylogeny*. Oxford University Press, Oxford, pp. 384–412.
- Rozanov, Y., 1969. *Probability Theory: A Concise Course*. Dover Publications.
- Saitou, N., Nei, M., 1986. The number of nucleotides required to determine the branching order of three species, with special reference to the human-chimpanzee-gorilla divergence. *J. Mol. Evol.* 24, 189–204.
- Steel, M., Székely, L., 2002. Inverting random functions II: explicit bounds for discrete maximum likelihood estimation, with applications. *SIAM. J. Discr. Math.* 15, 562–575.
- Townsend, J., 2007. Profiling phylogenetic informativeness. *Syst. Biol.* 56, 222–231.
- Townsend, J., Leuenberger, C., 2011. Taxon sampling and the optimal rates of evolution for phylogenetic inference. *Syst. Biol.* 60, 358–365.
- Townsend, J., Su, Z., Tekle, Y., 2012. Phylogenetic signal and noise: predicting the power of a data set to resolve phylogeny (in press). *Syst. Biol.* .